

Causal analysis, Correlation-Response, and Dynamic cavity

Erik Aurell

Dept Computational Biology, ACCESS Linnaeus Centre and Center for Quantum Materials, KTH-Royal Institute of Technology, SE-100 44 Stockholm, Sweden *and* Depts of Information and Computer Science and Applied Physics, Aalto University, P.O. Box 15400, FI-00076 Aalto, Finland

E-mail: eaurell@kth.se

Gino Del Ferraro

Dept Computational Biology, KTH-Royal Institute of Technology, SE-100 44 Stockholm, Sweden

E-mail: gino@kth.se

Abstract. The purpose of this note is to point out analogies between causal analysis in statistics and the correlation-response theory in statistical physics. It is further shown that for some systems the dynamic cavity offers a way to compute the stationary state of a non-equilibrium process effectively, which could then be taken an alternative starting point of causal analysis.

1. Causality in Philosophy, Physics and Statistics

Causality formalizes the universal human experience of agents (causes) taking actions leading to results (effects) ¹. In the Western Philosophical tradition Aristotle postulated four kinds of causes namely the material, the formal, the efficient and the final, out of which Bacon retained the material and the efficient. The material cause is in modern terms the material properties of an object, ultimately determined by which atoms it is composed of and how these are ordered in space, while the efficient cause is how the object is influenced by other objects. The Third Law of Newton [20] however states that in Nature there is no separation between cause and effect in the Aristotelian or Baconian sense; Physics fundamentally knows only interactions, and these are always mutual, a state of affairs unchanged since that time and the replacement of the Classical Physics by Quantum Physics. When the term causal is used to describe an interaction in modern high-energy Physics it means only that the influence cannot propagate faster than light so that object A at time t_A only depends on what happened at object B at times t_B early enough that a signal from B can reach A at time t_A , and *vice versa* [26].

The everyday and the philosophical notions of causality are in Physics instead intertwined with reversibility and irreversibility; Nature's laws are time-reversal invariant on the fundamental

¹ The literature on this topic is too vast and variegated to be meaningfully referenced; two classic studies from the perspective of belief systems in traditional societies can be found in [16] and [5].

level, but most ordinarily encountered processes are overwhelmingly likely to only flow in one direction [11, 7]. We say that dropping a glass vase on the floor is the cause of it breaking because it is exceedingly unlikely that the glass pieces would jump back together and fuse into a vase. Similarly, we say an enzyme causes a chemical reaction in one direction when the concentrations of the reactants are such that the opposite reaction is very unlikely. Although the details of this complex process are not fully known, we can also say that smoking causes cancer because the DNA in living cells is mostly that of one and the same genome for each individual – an extremely small subset of all possible DNA sequences of the same length – and cancerogenes in tobacco smoke therefore almost always lead to mutations away from the healthy genotype and into one out of very many deficient genotypes. In Nature cause-effect relationships are thus but abbreviations for processes in physical systems so strongly driven out of thermal equilibrium that they mostly only go one way.

Nevertheless, causal analysis is an important branch of statistics, describing the effects of interventions and answering questions of the “if-so-then-what?” character [21, 22]. Interventions are then taken to be outside Nature, typically ascribed to a human agent, and causality is thus distinct from statistical association studies. If person X is holding a glass vase and person Y trips him over, then person X is quite likely to fall and break the vase. However, we cannot know this for sure without observing the event as person X might for instance be much larger and stronger than Y . Likewise, if we can deactivate enzyme E then we can observe that a catalyzed reaction $S_1 \rightarrow S_2$ ceases, while if we can over-activate E then the catalyzed reaction goes faster. This is the paradigm for how molecular biologists identify interactions experimentally; good research practice and common sense hold that observing such direct responses is a more reliable means to acquire knowledge than observing the variations of E and the speed of the reaction $S_1 \rightarrow S_2$ in natural undisturbed conditions. For example, the catalyzed reaction may be one in a series of reactions $S_1 \rightarrow S_2 \rightarrow S_3 \rightarrow S_4 \rightarrow \dots$, catalyzed by enzymes E, E_2, E_3, \dots and the living cell may regulate the production of all these enzymes by the availability of the first substrate S_1 [1]. In this case E, E_2, E_3, \dots would all vary positively with the speed of the reaction $S_1 \rightarrow S_2$ and only a direct experiment can determine which of them actually does the job.

The first purpose of this paper is to show that there is a conceptual parallelism between causal analysis in statistics and long-time response functions in physics. A main difference is that the results of an intervention in causal analysis is (comparatively) easy to determine, while the response function integrates dynamical information on all time scales. The second purpose is to show how the tools of message-passing/Belief Propagation, developed to analyse complex static interdependencies, have been generalized to describe dynamics with complex interactions, and then known as dynamic cavity. Under suitable assumptions the dynamic cavity simplifies considerably the determination long-time responses, which can hence be considered an alternative starting point of a causal analysis. The paper is organized as follows. In Section 2 we describe in simple terms causal analysis, following mainly [23]. In Section 3 we give a short summary of correlation-response theory using Markov chains (synchronously updated spin systems) as our main example, and show the parallelism between a causal analysis and correlation-response. In Section 4 we describe dynamic cavity as a generalization of message-passing/Belief Propagation to dynamic phenomena, and show that it can be used to turn correlation-response into an alternative to causal analysis, for some systems. In Section 5 we sum up and discuss our results.

2. Causal analysis

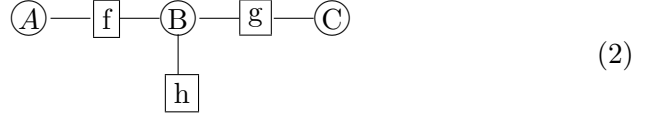
Statistical physicists are nowadays conversant with graphical models to describe probability distributions [27, 18]. To explain causal analysis we will start with two very simple Bayesian belief networks

$$(a) \quad \textcircled{A} \leftarrow \textcircled{B} \rightarrow \textcircled{C} \quad \text{and} \quad (b) \quad \textcircled{A} \rightarrow \textcircled{B} \rightarrow \textcircled{C} \quad (1)$$

where by $\textcircled{A} \rightarrow \textcircled{B}$ we mean that both the random variable B is dependent on random variable A in the ordinary sense of probability, and also that A (somehow) causes B . The dependency is encoded in conditional probabilities $P_{A|B}(a|b)$ where a and b are values of A and B . The joint probabilities of the three variables are in the two cases

$$\begin{aligned} \text{Case (a)} \quad P_B(b)P_{A|B}(a|b)P_{C|B}(c|b) &= \frac{P_{A,B}(a,b)P_{C,B}(c,b)}{P_B(b)} \\ \text{Case (b)} \quad P_A(a)P_{B|A}(b|a)P_{C|B}(c|b) &= \frac{P_{B,A}(ba)P_{C,B}(c,b)}{P_B(b)} \end{aligned}$$

where $P_B(b)$ is the marginal probability of variable B to take value b while $P_{A,B}(a,b)$ is the marginal probability of the pair of variables A and B to take values a and b , and so on. We assume for simplicity that $P(b)$ is different from zero for all values b of B . Both Bayesian belief networks encode the same joint probability; we cannot distinguish by covariation whether A causes B or B causes A . In the language of factor graphs this joint probability can alternatively be described by a factor graph [14]



with factors

$$f(a,b) = P(A=a, B=b) \quad g(c,b) = P(C=c, B=b) \quad h(b) = (P(B=b))^{-1} \quad (3)$$

and

$$P_{A,B,C}(a,b,c) = \frac{1}{Z} f(a,b)g(c,b)h(b) \quad (4)$$

The partition function Z is determined by the interactions encoded by the factors of the factor graph, and is in general difficult to compute for large models, but in the simple example considered here it is obviously equal to one.

In both cases described above the joint probability of A and C conditioned on B is then

$$P_{A,C|B}(a,c|\text{see}(B=b)) = \frac{P_{A,B}(a,b)P_{C,B}(c,b)}{(P_B(b))^2} = P_{A|B}(a|b)P_{C|B}(c|b) \quad (5)$$

where we have introduced Pearl's "see" notation [23]. To avoid confusion, let us note again that $P_{A|B}(a|b)$ in above has its ordinary probabilistic meaning of $\frac{\sum_c P_{A,B,C}(a,b,c)}{\sum_{a,c} P_{A,B,C}(a,b,c)}$ and is the same in both models.

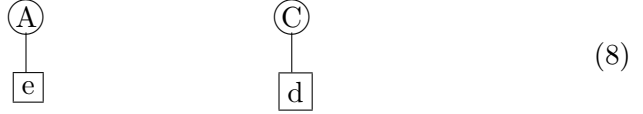
If we intervene on B and set its value to b the two Belief networks lead to new joint probabilities on the two remaining variables (A and C). In both cases A and C become independent with probabilities depending parametrically on the set value b , which we can call $P_A^{(B)}(a;b)$ ("the probability distribution of random variable A in the modified model where variable B has been set to constant b ") and $P_C^{(B)}(c;b)$ ("the probability distribution of random variable C in the modified model where variable B has been set to constant b "). Introducing Pearl's "do" notation [23] we then have

$$P_{A,C}(a,c|\text{do}(B=b)) = P_A^{(B)}(a;b)P_C^{(B)}(c;b) \quad (6)$$

and the dependencies can be illustrated as a (trivial) Bayesian network



or as an (equally trivial) factor graph



The two probability distributions are however not the same in the two cases.

In case (a) we have as numerical values $P_A^{(B)}(a; b) = P_{A|B}(a|b)$ and $P_C^{(B)}(c; b) = P_{C|B}(c|b)$, because A and C are then both assumed to be caused by B ; the factors in the factor graph (8) are $e = P_{A|B}(a|b)$ and $d = P_{C|B}(c|b)$. For this case (5) and (6) hence describe the same distribution.

In case (b) we also have $P_C^{(B)}(c; b) = P_{C|B}(c|b)$, but for the other probability instead $P_A^{(B)}(a; b) = P_A(a)$ corresponding to a factor $e = P_A(a)$ in (8). This difference is ultimately what it means to interpret the arrows in (1) as causes: if A is a cause and B is an effect then A should be unaffected by B , and in particular unaffected by any outside intervention on B . Therefore, whether or not there is any intervention on B , in case (b) the marginal probability of A is and remains $P_A(a)$ and the “do” (6) is different from the “see” (5). Expanding on the same point, in case (b) the “do-probability” $P_{A,C}(a, c|\text{do}(B = b))$ can be expressed in terms of probabilities observable before the intervention, namely as $P_A(a)P_{C|B}(c|b)$, but is not the same as the “see-probability” $P_{A,C}(a, c|\text{see}(B = b))$ which is, for both cases, $P_{A|B}(a|b)P_{C|B}(c|b)$. The Kullback-Leibler distance between the two is (for this simple example) $KL(\text{see}(B = b)|\text{do}(B = b)) = \sum_a P_{A|B}(a|b) \log \frac{P_{A|B}(a|b)}{P_A(a)}$, which generally is not zero.

We will now raise the abstraction level and following [23] define a general causal model M , also known as *Structural Equation Model*, as a set of exogenous variables U , a set of endogenous variables V_1, \dots, V_N , located in nodes $1, \dots, N$ in a graph G , for each node i a set of parent nodes $PA_i \subset \{1, \dots, N\} \setminus \{i\}$ and conditional probabilities $F_i(V_i|V_{PA_i}, U)$. The structure of G is determined by there being a link $i \rightarrow j$ iff $i \in PA_j$. Additionally one may include in the model specifications a distribution $P(U)$ over the exogenous variables [23]. Each such model defines a joint probability distribution of the endogenous variables as

$$P_M(V_1, \dots, V_N|U) = \frac{1}{Z_M(U)} \prod_i F_i(V_i|V_{PA_i}, U) \tag{9}$$

If G is a Directed Acyclic Graph (DAG), so that dependencies cannot propagate in a loop, clearly $Z_M(U) = 1$. The *do operator* is introduced by Pearl as:

Interventions and counterfactuals are defined through a mathematical operator called $do(x)$, which simulates physical interventions by deleting certain functions from the model, replacing them with a constant $X = x$, while keeping the rest of the model unchanged. The resulting model is denoted M_x . The postintervention distribution resulting from the action $do(X = x)$ is given by the equation

$$P_M(y|do(x)) = P_{M_x}(y)$$

Judea Pearl, “The Do-Calculus Revisited” (2012) [23]

It is useful to compare and contrast the do operation with the cavity method to be discussed in more detail below in Section 4. Both modify a probabilistic model by eliminating one or more variables, figuratively opening a hole (or cavity) in the factor graph. A first difference is that in the cavity method the variable and all its interactions are eliminated as if they were never there, while in a do operation the variable is set to a constant and the value of that constant

matters. Instead of (9) we thus have, taking $X = V_k$ for some k ,

$$P_{M_{v_k}} = \frac{1}{Z_{M_{v_k}}(U)} \prod_{\substack{i \neq k \\ k \notin V_{PA_i}}} F_i(V_i | V_{PA_i}, U) \prod_{\substack{i \neq k \\ k \in V_{PA_i}}} F_i^{V_k=v_k}(V_i | \{V_{PA_i}\} \setminus V_k, U) \quad (10)$$

where $Z_{M_{v_k}}$ is a new normalization constant and $F_i^{V_k=v_k}$ is a new function obtained from F_i by setting the variable V_k to the constant v_k . Hence we can express the “see” and the “do” as

$$P(V_i | \text{see}(V_k = v_k)) = \frac{P_M(V_i, V_k)}{P_M(V_k)} \quad P(V_i | \text{do}(V_k = v_k)) = P_{M_{v_k}}(V_i) \quad (11)$$

which again shows how and why the two concepts differ. A graphical illustration of the do operation given in Fig. 1.

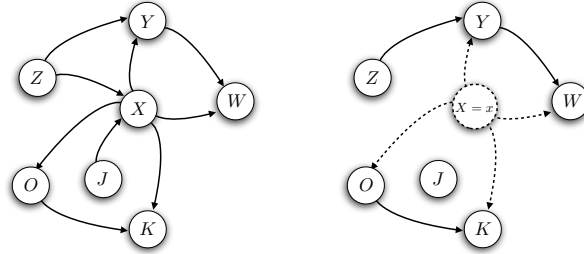


Figure 1. Illustration of the do operation. Left panel: a Bayesian belief network with a central node containing variable X . Right panel: reduced Bayesian belief network after intervening on variable X setting it to value x . Node containing X and outgoing links are indicated by dashed lines symbolizing that (10) depends parametrically on x . Ingoing links to node containing X are eliminated together with random variable X which does not appear in (10).

A second and more important difference is that the do operation is formulated for Bayesian belief networks of which (as we have seen) there can be many corresponding to the same joint probability distribution. Under the operations $\text{do}(X)$ for different X , each Bayesian belief network (*i.e.* each direction of the arrows in Eq. (1)) hence specifies a different set of changes of the joint probability distribution encoded in a factor graph.

An important question in causal analysis is whether probabilities after an intervention, *i.e.* $P_M(y|\text{do}(x))$, y standing for any subset of the endogenous variables, can be determined from observations before the intervention, *i.e.* from the set $P_M(z)$ (any and all z). When this is so one says that a causal effect query is *identifiable* because it can be decided (the probability $P_M(y|\text{do}(x))$ estimated) from data obtained before an intervention. In both the simple examples above this was the case, only the $P_M(y|\text{do}(x))$ ’s were not the same. More generally, a causal effect query is always identifiable from passively observed $P_M(z)$, provided that all variables in M are observed and G is known [22]. In less technical terms this last statement means nothing else than given sufficient data one can in principle estimate conditional probabilities, and given a direction of the arrows in a Bayesian belief network one can translate this information into what the conditional probabilities will be in a modified model; the situation is more complicated when some variables are unobserved (un-measured). The *Do-Calculus* of Pearl consists of three rules (a calculus, shown to be complete) for deciding identifiability when G is known and is a DAG, and some of the variables are unobserved. The Do-Calculus can hence be used to determine (in perhaps quite complex settings) whether a separate experiment is necessary, or

if an hypothetical question can be answered with the data already at hand. Other uses of Do-Calculus are *mediation*, how to separate direct and indirect causal relationships, *transportability*, how to determine how much information can be carried over to a new condition, and *meta-analysis*, how to pool data obtained under different conditions, for more details see [23].

3. Correlation-Response

At the basis of scientific mathematical philosophy is the idea that the regularities of the world are best expressed by how it changes in time, famously stated by Newton to Leibniz as *6accdae13eff7i319n4o4qrr4s8t12ux* [2]. Let us therefore substitute the Bayesian belief network in Section 2 by a minimal model encoding the same dependencies as a probabilistic evolution law:

$$P(V_1(t), \dots, V_N(t)) = \prod_{i=1}^N F_i(V_i(t) | V_{PA_i}(t-1)) P(V_1(t-1), \dots, V_N(t-1)) \quad (12)$$

The notation is here the same as in (9) except that the variables are now indexed by time (t) and a possible dependence on exogenous variables has been suppressed. Up to the technical simplification of synchronous dynamics, (12) is a prototype for a physically realistic mutual dependency. It could be realized in a biological regulatory system, say in a signal transduction network, where the cause-effect relationship between PA_i and V_i would have the underlying mechanistic interpretation of PA_i being the kinases, phosphatases and other enzymes catalyzing the phosphorylation, de-phosphorylation and other modifications to unit i . The endogenous variables U are then concentrations of molecules at constant concentrations, which could be sugars and other carbon sources for bacteria, or hormones and other signalling molecules in multi-cellular organisms. For long times the probability distribution in (12) would then reach stationary state which we will denote

$$P^*(V_1, \dots, V_N | U) = \lim_{t \rightarrow \infty} P(V_1(t), \dots, V_N(t) | U) \quad (13)$$

Since P^* in (12) is at least as realistic as P in (9) as a representation of how the endogenous variables depend on the exogenous variables we could also use it to define an analogy of the do operation. We thus have

$$P_M^*(y | \text{do}(x = X)) = P_{M_x}^*(y) = \lim_{\tau \rightarrow \infty} R(y, x, t + \tau, t) \quad (14)$$

where the last equality says that this is simply the long-time limit of a response function.

Generally a response function related to a generic quantity $V_i(t)$ is defined as $R_{ij}(t, t') = \langle \partial V_i(t) / \partial H_j(t') \rangle$ where $H_j(t')$ is a general parameter which can be varied within the system. An example, in ferromagnetic systems, is represented by the susceptibility function $\chi_{ij}(t, t') = \langle \partial M_i(t) / \partial H_j(t') \rangle$ which gives a measure of the change in the local magnetization $M_i(t)$ on site i at time t due to a change of an external field $H_j(t')$ acting on a different site j at an earlier time t' . In this respect, the *do* operation, otherwise stated as an intervention on the system, is similar to the variation of an (or several) external parameter(s) in a physical system which generates a response function. A main goal of causal analysis can be then reformulated as predicting suitable defined response functions under the change of some external tuneable parameters (intervention). A great deal is known about such response functions for systems at or near thermodynamic equilibrium where they are related to correlation functions through the Fluctuation-Dissipation-Theorem, which generically takes the form [15, 25]

$$\frac{1}{T} [C(\tau = 0) - C(\tau)] = \int_0^\tau R(\tau') d\tau' \quad (15)$$

where $\tau = t - t'$, time translational invariance is assumed, and T is temperature. We note that the left hand side of above is measured in the unperturbed system and the right hand side in the perturbed system. Causal effect queries are therefore always identifiable in systems at or near thermodynamic equilibrium from observing no more than the correlation between the variable which is set and the variable one wants to predict. The relation between correlation and response has been also used to improve for inference on a network [24].

4. Dynamic cavity

In this section we describe how the techniques now generally called message-passing or Belief Propagation can be generalized to analyse evolution laws like (12). Message-passing techniques have been invented independently in different fields [27, 14]. In Physics they are also known as the cavity method [18], and usually traced back to [3]. Their purpose is to compute marginal probabilities over some (usually small) subset of variables in a probabilistic model described by a factor graph which is done by storing partial computations in nodes representing the variables and then forwarding such partial results to neighbours in the graph for further processing. Message-passing converges and is exact if the underlying graph is a tree but also often converges and is a very good approximation if the underlying graph has only long loops, a fact that has many theoretical and practical applications in coding theory and elsewhere [18]. The fixed points of the algorithms correspond to stationary points under variation of the Bethe approximation to the free energy in the corresponding statistical mechanics problem [27, 18].

Our point of departure is now the observation that the dynamics (12) naturally leads to a probability distribution on variable histories

$$P(X_1, \dots, X_N) = P(V_1(0), \dots, V_N(0)) \prod_{t=1}^T \prod_{i=1}^N F_i(V_i(t) | V_{PA_i}(t-1)) \quad (16)$$

where $X_i = \{V_i(0), \dots, V_i(T)\}$ for $i = 1, \dots, N$. Before continuing, let us note that if the variables are Boolean and take values $\{-1, 1\}$ (“spins”) then (12) specifies a dynamics of a spin system under synchronous updates, and if further all the transition probabilities are of the type $F_i(V_i(t) | V_{PA_i}(t-1)) \propto \exp\left(V_i(t) \left(h_i + \sum_j J_{ij} V_j(t-1)\right)\right)$ is known as the Kinetic Ising model [12]. The parent set V_{PA_i} is then comprised of the variables V_j for which J_{ij} is non-zero. When $J_{ij} = J_{ji}$ for all pairs (i, j) the system has a stationary state $P(V_1, \dots, V_N) \propto \exp\left(\sum_i h_i V_i + \sum_{ij} J_{ij} V_i V_j\right)$ and (12) then simulates a system in thermal equilibrium, albeit under the somewhat unphysical synchronous update rule. In the more general case when $J_{ij} \neq J_{ji}$, and in particular for fully asymmetric models where J_{ij} can only be non-zero when J_{ji} equals to zero, (12) on the other hand simulates a non-equilibrium system.

The first result on reducing the complexity of (16) dates back almost thirty years [8] and pertains to fully asymmetric models. For these an influence $X_j \rightarrow X_i$ must traverse a loop in G to get back to X_j , and when there are no loops, or when these can otherwise be disregarded, the marginal probability of X_j is independent of X_i . This leads to simple equation for the marginalization over a single variable and and single time, namely

$$P_i(V_i, t) = \sum_{V_j \in PA_i} F_i(V_i | V_{PA_i}) \prod_j P_j(V_j, t-1) \quad (\text{Fully asymmetric}) \quad (17)$$

We now generalize a bit and assume that the dependency graph G has the property associated with the effectiveness of standard message-passing *i.e.* that it is a tree, or at least locally tree-like. That is, we assume that one cannot form circular dependency chains $i \rightarrow j \rightarrow k \rightarrow \dots \rightarrow i$ where $V_i \in PA_j, V_j \in PA_k, \dots \in PA_i$ unless either somewhere the chain backtracks as

$\dots j \rightarrow k \rightarrow j \rightarrow \dots$ or the chain is long, on the order of the graph diameter of G . The term “dynamic cavity” was introduced in [13] for such situations where it was used to obtain rigorous bounds on the consensus threshold for the majority dynamics. Methods have later been developed to treat, in principle exactly, such problems when the dynamical law is modified to only allow transitions in one direction [17, 6]. An important step was taken in [19] where marginals in a stationary state were computed approximately based on an ansatz, recently extended to to a perturbative scheme and to also cover transient phenomena [10, 4]. The main problem is then that even when the dependency graph G of the dynamics in (12) is locally tree-like this is not the case for dependencies in (16) due to “loops-in-time”. These dependencies have been resolved by a graph expansion technique [17, 6, 10] which we now explain.

First, as for the Kinetic Ising model it is often convenient to define transition functions only up to a normalization $F_i(V_i|V_{PA_i}) \propto \exp(r_i(V_i, V_{PA_i}))$. The normalization constant is then $N_i(V_{PA_i}) = \sum_{V_i} \exp(r_i(V_i, V_{PA_i}))$, a function that does not depend explicitly on V_i . Assuming further for simplicity that interaction functions r_i are only pair-wise the dependency graph can be illustrated as in Fig. 2. The model defined on variable histories, (16), now has short-loop dependencies even if the graph G itself does not. This can be seen by tracing the dependency of one of these variable, say $X_i = \{V_i(0), \dots, V_i(T)\}$. Pick a time t and notice that $V_i(t)$ depends on $V_j(t-1)$ for all $j \in PA_i$. Then pick two of these variables X_j and X_k such that $i \in PA_j$ and $i \in PA_k$, then $V_j(t-1)$ and $V_k(t-1)$ both depend on $V_i(t-2)$. At the same time $V_j(t-1)$ and $V_k(t-1)$ are however also dependent through the normalization $N_i(V_{PA_i})$, and X_i , X_j and X_k are therefore connected in a dependency loop of length three.

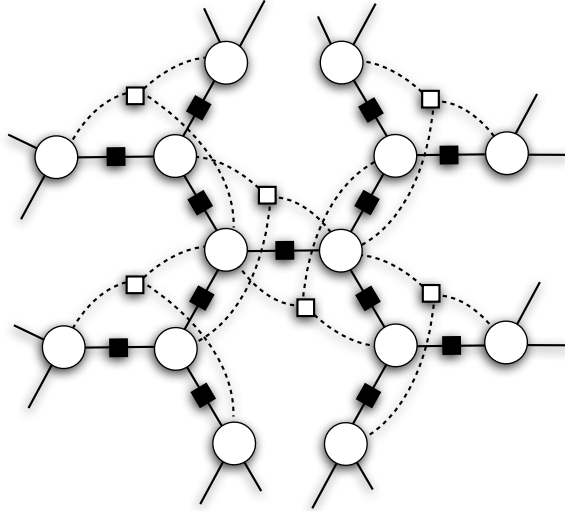


Figure 2. A tree-like dependency graph with the normalization constants in the transition functions split off as separate factor nodes (boxes). It has been assumed that the dependencies are not fully asymmetric so that when node i depends on node j , node j in general also depends on node i . Dependencies between nodes (in general mutual) are indicated by (undirected) lines. In the kinetic model loops emerge from variables at different times in (16).

While there are many approaches to get rid of loops in factor graphs we will use one which is well adapted to the dynamics. For every pair i and j such that $V_i \in PA_j$ (whether or not also $V_j \in PA_i$) we introduce a new compound variable $(X_i^{(ij)}, X_j^{(ij)})$ interpreted as “variable $X_i^{(ij)}$ of type X_i belonging to link (i, j) and $X_j^{(ij)}$ of type X_j also belonging to link (i, j) ”². Introducing

² To be precise the two parts of the compound variable are distinguished by their index (i or j) and not by

now the consistency requirement that the variables $X_i^{(ij)}$ take the same value for all the links (i, j) where this type of variable is found we can rewrite (16) as

$$\begin{aligned} P(\{X_i^{(ij)}, X_j^{(ij)}\}) &= P_{\text{init}} \prod_{t=1}^T \prod_i F_i(\overline{V_i(t)} | \{V_j^{(ij)}(t-1)\}_{j \in PA_i}) \\ &= \prod_{t=1}^T \prod_i \mathbf{1}_{V_i^{(ij_1)}(t)=V_i^{(ij_2)}(t)=\dots} \end{aligned} \quad (18)$$

where P_{init} is the probability distribution on the initial conditions translated to the new variables, $V_j^{(ij)}(t)$ are the restrictions of the variables $X_j^{(ij)}$ to a single time t and $\overline{V_i(t)}$ is any suitable average of the $V_i^{(ij)}(t)$ for different j [10]. The graph expansion is illustrated in Fig. 3. Introducing messages in the standard way and summing out the consistency conditions we thus

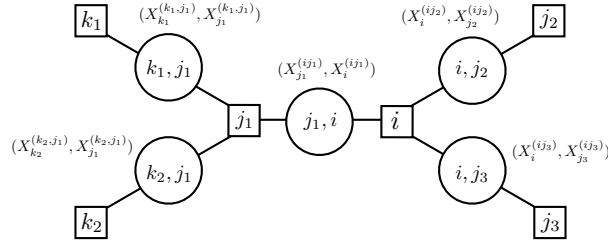


Figure 3. Auxiliary graph obtained from a part of Fig. 2 where loops have been removed by a graph expansion procedure. The new variable nodes contain histories of two variables that were neighbours in the original graph while new factor nodes (one per each old variable node) contain both the transition functions and the consistency conditions.

arrive at [10]

$$m_{i \rightarrow (ij)}(X_i^{(ij)}, X_j^{(ij)}) \propto \sum_{\{X_k^{(ik)}\}} \Phi_i(X_i^{(ij)}, X_j^{(ij)}, \{X_k^{(ik)}\}) \prod_{k \in \partial i \setminus j} m_{k \rightarrow (ik)}(X_k^{(ik)}, X_i^{(ij)}) \quad (19)$$

where $\Phi_i(X_i, X_j, \{X_k\}) = \prod_{t=1}^T F_i(V_i(t) | V_j(t-1), \{V_k(t-1)\}_{k \in PA_i \setminus j})$. Equation (19) are the dynamic cavity update equations which are the same as ordinary cavity update equations applied to the model (16) on variable histories. A trace of the dynamic origin remains in that the probability $m_{i \rightarrow (ij)}(X_i^{(ij)}, X_j^{(ij)})$ can be taken to depend on the full history $X_i^{(ij)} = \{V_i^{(ij)}(0), \dots, V_i^{(ij)}(T)\}$ of the first argument, but only a one unit shorter history $X_j^{(ij)'} = \{V_j^{(ij)}(0), \dots, V_j^{(ij)}(T-1)\}$ of the second argument. For a discussion as well as a description of the analogous dynamic cavity output equations, see [10]. To make (19) practical further assumptions are needed, to close the iterations in a low-dimensional subspace of the functions $m_{i \rightarrow (ij)}(X_i^{(ij)}, X_j^{(ij)})$. In [10] good results were reported based on closure in the class of 1-step Markov processes, leading to schemes not much more complicated than (17) while in [4] even better results were reported from more involved procedure. The field is in active development and likely even better approximations will appear in the near future.

We will now take the point of view that the probabilities P^* and $P_{M_x}^*$ in (13) and (14) are efficiently computable and ask what are the implications for causal analysis. First, the

their order in the pair. When a message is to be transmitted from i to j they are naturally read in the order $(X_i^{(ij)}, X_j^{(ij)})$ while if the message is transmitted in the opposite direction the natural order is $(X_j^{(ij)}, X_i^{(ij)})$.

assumption of synchronous updates is unrealistic in most natural systems but certainly no more so than the assumption of instantaneous dependence made in (9). In most problems where an underlying mechanistic explanation is conceivable “causes” are ultimately to be interpreted as variables influencing transition rates, and the simplest example of such dynamics is (12). In stationary state an underlying explanation, which one could call “mechanistic causes”, leads to a joint probability distribution P^* with generally many more dependencies. That is, there will be one (directed) dependency graph G describing the probabilistic evolution law (12) and another (undirected) factor graph F describing the probability P^* in (13), and F will almost always be (much) larger and (much) richer than G . For a worked-out example of such an effect, in the relaxation towards equilibrium of the Kinetic Ising model on a 1D lattice [12], see [9].

The stark conclusion of these in essence elementary considerations is hence simply that causal analysis should be treated with a great deal of caution. The causes identified by causal analysis may not be causes at all in the everyday understanding of the word. Instead they may incorporate all kinds of indirect effects from forcing a stationary state of a dynamics to be represented by a Bayesian belief network, potentially masking a far simpler underlying reality.

5. Summary and discussion

We have given a brief introduction to causal analysis and discussed how it extends the tools of factor graphs and probabilistic models to describe outside interventions that change the models themselves. We have compared and contrasted causal analysis to the analysis of dynamic processes by the dynamic cavity method and pointed out similarities between causal analysis and physical correlation-response theory.

The *do calculus* considers causal relationships to be the fundamental building blocks of reality [22] and aims to discover which are these casual dependences in the system under investigation (usually DAG networks). Although this certainly is a fascinating goal, it is at odds with physical theory which does not admit causes and effects in the philosophical sense on the fundamental level, but only for macroscopic (irreversible) processes. For such processes the flow of time is however essential, and causes are thus naturally understood as variables influencing transition rates between various states in a system. The stationary states of such processes are normally quite complicated reflecting not only the dependencies in the transition rates, but also chains of such dependencies of arbitrary length, the only major exception being systems in thermodynamic equilibrium. Therefore, caution is needed when interpreting the results of causal analysis as causes in an everyday – *e.g.* legal – sense. The mechanisms identified by causal analysis include (except in thermal equilibrium) both underlying direct effects and many kinds of indirect effects where the setting of one variable influences the behaviour of another at a later time through one or many intermediaries.

The major advantage of causal analysis is instead in its relative simplicity of its basic ansatz. Up to recent times few methods except Monte Carlo simulations were available to analyse the dynamics of non-equilibrium systems, and determining their stationary states was therefore laborious. We have discussed that for some systems the dynamic cavity offers an alternative approach which could therefore be used as an alternative starting point of causal analysis. Many major issues however remain to be solved in that approach, the most important one perhaps being how to extend the dynamic cavity (if possible) to continuous-time processes.

Acknowledgments

E.A. acknowledges valuable discussions with Dr. Antti Hyttinen and Profs. Timo Koski and Jukka Corander. This work has been supported by the European Union through Marie Curie ITN “NETADIS” (GDF), by Swedish Science Council through grant 621-2012-2982 (EA), and by the Academy of Finland through its Center of Excellence COIN (EA).

References

- [1] Uri Alon, “An introduction to Systems Biology”, Chapman & Hall/CRC (2007)
- [2] Isaac Newton, Letter to Leibniz on October 24 1676, *Opera omnia IV*, 540 ff, to be rendered as “Data æquatione quocunque fluentes quantitatæ involvente fluxiones invenire et vice versa” [Given in an equation the fluents of any number of quantities, to find the fluxions and vice versa]. Quoted by V.I. Arnold, “Geometrical Methods in the Theory of Ordinary Differential Equations”, Preface to First Edition, Springer (1983:1989) and there translated as “It is useful to solve differential equations”.
- [3] Hans A. Bethe, “Statistical theory of superlattices”, *Proc. Roy. Soc. London A* **150** 552558 (1935)
- [4] Thomas Barthel, Caterina De Bacco, Silvio Franz “A matrix product algorithm for stochastic dynamics on locally tree-like graphs”, [arXiv:1508.03295] (2015)
- [5] Pierre Bourdieu, “Le sens pratique”, Les Éditions de Minuit, Collection “Le sens commun”, Paris (1980)
- [6] F. Altarelli, Alfredo Braunstein, Luca Dall’Asta, A. Ingrosso, Riccardo Zecchina, “The zero-patient problem with noisy observations”, *J. Stat. Mech.* P10016 (2014)
- [7] Sergio Chibbaro, Lamberto Rondoni, Angelo Vulpiani “Reductionism, Emergence and Levels of Reality”, [Chapter 4: From Microscopic Reversibility to Macroscopic Irreversibility] Springer (2014)
- [8] Bernard Derrida, Elizabeth Gardner, Anne Zippelius, “An exactly solvable asymmetric neural network model” *EPL* **4** 167 (1987)
- [9] Gino Del Ferraro, Erik Aurell “Perturbative large deviation analysis of non-equilibrium dynamics”, *J. Phys. Soc. Japan* **83** 084001 (2014)
- [10] Gino Del Ferraro, Erik Aurell, “Dynamic message-passing approach for kinetic spin models with reversible dynamics”, *Phys. Rev. E* **92**, 010102 (2015)
- [11] Richard P. Feynman, “The Character of Physical Law”, [Lecture 5: The Distinction of Past and Future] The MIT Press (1964)
- [12] Roy J. Glauber, “Time-Dependent Statistics of the Ising Model” *J. Math. Phys.* **4** 294 (1963)
- [13] Yashodhan Kanoria, Andrea Montanari, “Majority dynamics on trees and the dynamic cavity method”, *The Annals of Applied Probability* **21** 1694–1748 (2011)
- [14] Frank R. Kschischang, Brendan J. Frey, Hans-Andrea Loeliger “Factor Graphs and the Sum-Product Algorithm”, *IEEE transactions on Information Theory* **47** 498-519 (2001)
- [15] Kubo, R. (1966). The fluctuation-dissipation theorem. *Reports on progress in physics*, 29(1), 255.
- [16] Claude Lévi-Strauss “La pensée sauvage”, Plon, Paris (1962)
- [17] Andrey Y Lokhov, Marc Mézard, Lenka Zdeborová, “Dynamic message-passing equations for models with unidirectional dynamics”, *Physical Review E* **91** (1), 012811 (2015)
- [18] Marc Mézard, Andrea Montanari, “Information, physics, and computation”, Oxford University Press (2009)
- [19] Isaak Neri, Desiré Bollé, “The cavity approach to parallel dynamics of Ising spins on a graph” *J. Stat. Mech.* P08009 (2009)
- [20] Isaac Newton, *Principia* (1687) “Lex III: Actioni contrariam semper et æqualem esse reactionem: sive corporum duorum actiones in se mutuo semper esse æquales et in partes contrarias dirigi” [Law III: To every action there is always opposed an equal reaction: or the mutual actions of two bodies upon each other are always equal, and directed to contrary parts]
- [21] Judea Pearl “Causal inference in statistics: An overview”, *Statistics Surveys* **3** (2009) 96146
- [22] Judea Pearl *Causality* (2000, 2nd edition 2009) Cambridge University Press
- [23] Judea Pearl, “The Do-Calculus Revisited”, Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence (UAI2012) August 1517, 2012 Catalina Island US, p.4 [arXiv:1210.4852]
- [24] Raymond, J. and Ricci-Tersenghi F. (2013, June). Correcting beliefs in the mean-field and Bethe approximations using linear response. In *Communications Workshops (ICC), 2013 IEEE International Conference on (pp. 1429-1433). IEEE*.
- [25] Villamaina D., Baldassarri A., Puglisi A. and Vulpiani A. (2009). The fluctuation-dissipation relation: how does one compare correlation functions and responses? *Journal of Statistical Mechanics: Theory and Experiment*, 2009(07), P07024.
- [26] Steven Weinberg *The Quantum Theory of Fields: volume 1 Foundations*, Cambridge University Press, Cambridge UK (1995)
- [27] Jonathan S. Yedidia, William T Freeman, Yair Weiss, “Understanding belief propagation and its generalizations”, *Exploring artificial intelligence in the new millennium*, **8**, 236–239 (2003)